

Marginally Interpretable Generalized Linear Mixed Models

Jeffrey J. Gory, Peter F. Craigmile, and Steven N. MacEachern

Department of Statistics
The Ohio State University
Columbus, OH 43210

Abstract

Two popular approaches for relating correlated measurements of a non-Gaussian response variable to a set of predictors are to fit a *marginal model* using generalized estimating equations and to fit a *generalized linear mixed model* by introducing latent random variables. The first approach is effective for parameter estimation, but leaves one without a formal model for the data with which to assess quality of fit or make predictions for future observations. The second approach overcomes the deficiencies of the first, but leads to parameter estimates that must be interpreted conditional on the latent variables. Further complicating matters, obtaining marginal summaries from a generalized linear mixed model often requires evaluation of an analytically intractable integral or use of attenuation factors that are not exact. We define a class of marginally interpretable generalized linear mixed models that lead to parameter estimates with a marginal interpretation while maintaining the desirable statistical properties of a conditionally-specified model. We discuss the form of these models under various common link functions and also address computational issues associated with these models. In particular, we introduce an accurate and efficient method for evaluating the logistic-normal integral that commonly arises with binary response data.

Keywords: logistic-normal integral, marginal model, conditional model, population-averaged, subject-specific predictions

1 Introduction

The generalized linear mixed model (GLMM) combines the generalized linear model (see Nelder and Wedderburn, 1972; McCullagh and Nelder, 1989) with the linear mixed model (see Henderson et al., 1959; Henderson, 1975; Laird and Ware, 1982) to form a model that allows for non-Gaussian responses as well as random effects. Such a model typically relates a linear combination of p predictors \mathbf{x} and q random effects \mathbf{U} to a response Y through a *link function* $g(\cdot)$ (see Stiratelli et al., 1984; Gilmour et al., 1985; McCulloch et al., 2008). It has the form

$$\mu = E(Y|\mathbf{U} = \mathbf{u}) = h(\mathbf{x}^T \boldsymbol{\beta} + \mathbf{d}^T \mathbf{u}), \quad (1)$$

where $\boldsymbol{\beta}$ is a p -vector of fixed effect parameters, \mathbf{d} is a q -vector of covariates, and $h(\cdot) = g^{-1}(\cdot)$ is the inverse link function. To complete the specification of the model one typically assumes a distributional form for the random effects \mathbf{U} and the response Y given \mathbf{U} . This defines a *conditional model* for which the mean μ is conditioned on the random effects \mathbf{U} .

Interest often lies in *marginal* or *population-averaged* effects rather than conditional effects. Although one could obtain marginal predictions from a GLMM, it is common to directly model the marginal mean using what is known as a *marginal model*. Such a model involves specification of a mean structure, typically written as $\mu = E(Y) = h(\mathbf{x}^T \boldsymbol{\beta})$, and a covariance structure, usually with no distributional form explicitly assumed for the data. Estimation of $\boldsymbol{\beta}$ for this model, ordinarily accomplished using *generalized estimating equations* (GEE), can be robust to misspecification of the covariance structure (Liang and Zeger, 1986; Zeger and Liang, 1986). However, the absence of a likelihood function in the marginal model makes it difficult to check and compare models. Further, since the GEE method does not assume a generative model, one cannot easily make predictions at the individual level.

The distinction between marginal and conditional models is important because the parameters in these two types of models are generally not the same when the link function $g(\cdot)$ is nonlinear (see Zeger et al., 1988; Neuhaus et al., 1991; Neuhaus and Jewell, 1993; Diggle et al., 2002; Ritz

and Spiegelman, 2004). Several authors have addressed this discrepancy by attempting to find a proportional relationship between the marginal and conditional parameters. For example, Wang and Louis (2003) proposed a *bridge distribution* f_U for univariate random effects, which depends on the inverse link function $h(\cdot) = g^{-1}(\cdot)$, such that for some constant c and all $\mathbf{x}^T \boldsymbol{\beta} \in \mathbb{R}$

$$\int h(\mathbf{x}^T \boldsymbol{\beta} + u) f_U(u) du = h(c \mathbf{x}^T \boldsymbol{\beta}). \quad (2)$$

Similarly, Caffo et al. (2007) proposed altering the inverse link function $h(\cdot)$ in order to satisfy (2). Earlier work by Zeger et al. (1988) and Neuhaus et al. (1991) involved finding *attenuation factors* a ($0 < a < 1$) such that (2) holds approximately for a range of $\mathbf{x}^T \boldsymbol{\beta}$ with $c = 1/a$.

Although these strategies recognize the difference between marginal and conditional models, they fail to provide a single model that both has parameters with a marginal interpretation and allows one to easily make predictions at the individual level. Instead, they provide a relationship that one could use to obtain parameters with an alternative interpretation after a conditional or marginal model has been fit. In this paper, we explore a class of conditionally-specified, likelihood-based models with a direct marginal interpretation for the parameters. We say that a GLMM is *marginally interpretable* if and only if for all $\mathbf{x}^T \boldsymbol{\beta}$

$$\int h(\mathbf{x}^T \boldsymbol{\beta} + \mathbf{d}^T \mathbf{u}) f_{\mathbf{U}}(\mathbf{u}) d\mathbf{u} = h(\mathbf{x}^T \boldsymbol{\beta}), \quad (3)$$

where $f_{\mathbf{U}}$ is the joint density of the random effects. This property does not hold for most common choices of link function $g(\cdot)$ and random effects density $f_{\mathbf{U}}$. In particular, when mean-zero normal random effects are paired with canonical link functions, (3) is generally not satisfied.

The *marginalized multilevel model* (Heagerty, 1999; Heagerty and Zeger, 2000) satisfies (3). This model involves reparameterizing the conditional mean model as $E(Y|\mathbf{U}) = h(\Delta + \mathbf{d}^T \mathbf{U})$, where Δ is a function of $\mathbf{x}^T \boldsymbol{\beta}$ and $f_{\mathbf{U}}$ that is defined implicitly by the equation

$$\int h(\Delta + \mathbf{d}^T \mathbf{u}) f_{\mathbf{U}}(\mathbf{u}) d\mathbf{u} = h(\mathbf{x}^T \boldsymbol{\beta}). \quad (4)$$

An advantage of the marginalized multilevel model is that it can be more robust to misspecification of the random effects distribution than a conventional GLMM (Heagerty and Kurland, 2001).

In this article we parameterize $\Delta = \mathbf{x}^T \boldsymbol{\beta} + \mathbf{d}^T \mathbf{a}$, where $\mathbf{d}^T \mathbf{a}$ is known as the *adjustment* and is defined such that (4) holds. The adjustment $\mathbf{d}^T \mathbf{a}$ potentially depends on the fixed portion of the model $\mathbf{x}^T \boldsymbol{\beta}$, the parameters characterizing the random effects distribution f_U , and the random effects design \mathbf{d} . Despite being a deterministic piece of the model, $\mathbf{d}^T \mathbf{a}$ is viewed as a location shift of the random effects distribution. As such, we cease to conceptualize each random effect as a uniform shift in the response that is associated with a group or cluster of observations because the magnitude of the shift for a specific observation within a group or cluster depends on the covariate values for that particular observation. An example of when different observations with the same random effect could be shifted by different amounts is when there are repeated measurements on an individual over time and the covariates vary with time. Additional details regarding this characterization of the random effects can be found in Section 2.2.

In general, parameter estimation for GLMMs can be difficult because evaluation of the marginal likelihood, which is required for maximum likelihood estimation, often involves an analytically intractable integral. Thus, one must use numerical integration or employ a method that approximates the likelihood, such as *penalized quasi-likelihood* (PQL) (Breslow and Clayton, 1993). An alternative is to adopt a Bayesian framework and to employ Markov chain Monte Carlo (MCMC) to produce samples from the posterior for inference (e.g. Zeger and Karim, 1991; Gamerman, 1997; Miglioretti and Heagerty, 2004). The adjustment in our proposed model can easily be incorporated into an MCMC algorithm and, although not always available in closed form, can be calculated efficiently. In particular, we will discuss a novel approach for evaluating the logistic-normal integral that allows for accurate and efficient calculation of the adjustment for a model with a logit link and normal random effects.

In the next section we formally introduce our proposed characterization of the model and discuss its properties. Section 3 provides more detail about the form of the model for specific link

functions. In Section 4 we discuss a Bayesian approach to fitting the proposed model, including an algorithm for efficiently computing the logistic-normal integral. Applications of the proposed model to data from a teratological experiment on rats and from a clinical trial on epileptics are given in Section 5. Finally, Section 6 contains a discussion of possible avenues for further research. Proofs of all results are provided in the Appendix. Additional details about how to fit the model can be found in the Supplementary Material.

2 Marginally Interpretable Generalized Linear Mixed Models

Formally, we propose the following mixed effects model. For a response Y_i ($i = 1, \dots, N$), a p -vector of predictors \mathbf{x}_i , and a q -vector of random effects \mathbf{U}_i , we model the conditional mean as

$$E(Y_i | \mathbf{U}_i = \mathbf{u}) = h(\mathbf{x}_i^T \boldsymbol{\beta} + \mathbf{d}_i^T \mathbf{u} + \mathbf{d}_i^T \mathbf{a}_i), \quad i = 1, \dots, N, \quad (5)$$

where the \mathbf{U}_i have joint density given by $f_{\mathbf{U}}(\mathbf{U}_i)$ and, conditional on the \mathbf{U}_i , the Y_i are mutually independent with density $f_{Y|\mathbf{U}}(Y_i | \mathbf{U}_i)$ for each i . The adjustment $\mathbf{d}_i^T \mathbf{a}_i$, when it exists, is defined implicitly by the equation

$$h(\mathbf{x}_i^T \boldsymbol{\beta}) = \int h(\mathbf{x}_i^T \boldsymbol{\beta} + \mathbf{d}_i^T \mathbf{u} + \mathbf{d}_i^T \mathbf{a}_i) f_{\mathbf{U}}(\mathbf{u}) d\mathbf{u}, \quad (6)$$

and is included to ensure that the model we consider is marginally interpretable as defined in (3).

By including the adjustment in our GLMM, we have specified a formal statistical model for which the parameters $\boldsymbol{\beta}$ have a marginal interpretation. Likelihood-based methods can be employed to fit this model and the results can be used to make inferences about the marginal mean and to make predictions at the individual level. The proposed model is superior to a typical marginal model defined only in terms of the mean and covariance structure because it is a fully-specified model with a likelihood. Further, we argue that the proposed model could be preferred to a conventional GLMM that does not include the adjustment because it provides a direct marginal interpretation of the parameters.

2.1 Replication and Consistency

The random effects in a mixed model provide a means of introducing dependence and overdispersion into the model. For example, repeated measures on a subject may be systematically large, or count data may show extra-Poisson variation. In simple models, the degree of dependence/overdispersion is determined by the distribution of the random effects, most typically by its variance.

The twin notions of dependence and overdispersion lead to the tenet that one must have replication to consistently estimate the random effects variance σ^2 . As confirmation of this point, consider the following hierarchical model:

$$\begin{aligned} Z_i &\sim F_{\sigma^2}; \\ Y_i|Z_i = z_i &\sim \text{Bernoulli}(z_i), \end{aligned}$$

where $i = 1, \dots, N$ and F_{σ^2} is an arbitrary distribution on $(0, 1)$ with mean μ and variance σ^2 . The marginal model is found by integrating over Z_i and the resulting distribution of Y_i is $\text{Bernoulli}(\mu)$, irrespective of σ^2 . No matter how many of these Bernoullis are collected, there is no replication tied to a single random effect, *no* information is obtained about σ^2 , and hence σ^2 cannot be estimated consistently.

Against this background, Kim and Kim (2011) proved a striking result for a conventional Bernoulli GLMM—namely that the maximum likelihood estimator is strongly consistent for the random effects variance σ^2 . We call this the Kim Paradox, stating a slightly different result than is presented in Kim and Kim (2011) and placing the result in our notation.

The Kim Paradox: *With no replication, one can estimate σ^2 consistently.* Let parameters β_0 , $\beta_1 \neq 0$, and $\tau^2 > 0$ be fixed and known, and let $X_i \sim N(0, \tau^2)$ and $U_i \sim \text{Uniform}(-c, c)$ ($c > 0$, $i = 1, 2, \dots$), be independent sequences of random variables. Furthermore, let $g(\cdot) = \text{logit}(\cdot)$, and define the conditionally independent sequence $Y_i|X_i = x_i, U_i = u_i \sim \text{Bernoulli}\{h(\beta_0 + \beta_1 x_i + u_i)\}$ for $i = 1, 2, \dots$. Then $\hat{\sigma}^2$, the maximum likelihood estimator of σ^2 , is consistent.

The Kim Paradox lives at the intersection of marginal and conditional GLMMs. It arises from the fact that $E\{h(\beta_0 + \beta_1 x_i + U_i)\}$, the expectation of the conditional mean in a conventional GLMM, is distorted by the random effects in such fashion that there is a 1 – 1 mapping between σ^2 and $E\{h(\beta_0 + \beta_1 x_i + U_i)\}$. This 1 – 1 mapping, along with a rich enough set of x_i , ensures that the marginal mean functions are identifiable and consistency of σ^2 follows.

A marginally interpretable Bernoulli GLMM of the form (5) and (6) resolves the Kim Paradox. The marginal mean is unaffected by changes in σ^2 . Consequently, the data contain *no* information about σ^2 , and one cannot obtain a consistent estimator of the random effects variance without replication. This is stated more formally in the following proposition:

Proposition 1. *If $Y_i|U_i$ is Bernoulli-distributed and we have a marginally interpretable GLMM of the form given by (5) and (6) for which the random intercepts U_i , $i = 1, \dots, N$, are independently distributed, then the marginal density of Y_i does not depend in any way on the distribution of U_i .*

2.2 Interpretation of Random Effects

Traditionally, one views each random effect as a shift in the mean response that applies to all of the units in a group of observations sharing that random effect. As such, random effects account for additional variation across groups and help to model situations where observations within the same group are generally more similar than observations in different groups. When the curvature of the link function is not uniform across the range of the covariates it may be necessary to shift each unit within the same group of observations by a different amount in order to preserve the marginal mean. Thus, the idea of shifting all units sharing the same random effect by the same amount is not always applicable in a marginally interpretable model.

The role of the adjustment in the proposed model is to ensure that (3) holds by accounting for the curvature of the inverse link function. The form of the adjustment is determined by the choice of link function and random effects distribution, whereas its specific value typically depends on $\mathbf{x}_i^T \boldsymbol{\beta}$. Exceptions for which $\mathbf{d}_i^T \mathbf{a}_i$ does not depend on $\mathbf{x}_i^T \boldsymbol{\beta}$ are models with an identity link or a log

Table 1: Form and existence of the adjustment for common link functions

Link Function	Distribution of U_i	Form of $\mathbf{d}_i^T \mathbf{a}_i$	Closed Form?
identity	mean exists and equals zero	zero	yes
log	exponential tails	independent of $\mathbf{x}_i^T \boldsymbol{\beta}$	yes
probit	Gaussian	linear in $\mathbf{x}_i^T \boldsymbol{\beta}$	yes
	non-Gaussian	nonlinear in $\mathbf{x}_i^T \boldsymbol{\beta}$	no
logit	bridge distribution	linear in $\mathbf{x}_i^T \boldsymbol{\beta}$	yes
	most other distributions	nonlinear in $\mathbf{x}_i^T \boldsymbol{\beta}$	no
complementary log-log	bridge distribution	linear in $\mathbf{x}_i^T \boldsymbol{\beta}$	yes
	most other distributions	nonlinear in $\mathbf{x}_i^T \boldsymbol{\beta}$	no
square root	restrictions on domain	nonlinear in $\mathbf{x}_i^T \boldsymbol{\beta}$	yes
reciprocal	$E\{1/(\mathbf{x}_i^T \boldsymbol{\beta} + \mathbf{d}_i^T \mathbf{U}_i)\}$ exists	see Section 3.3	see Section 3.3

link. For a model with an identity link, (3) holds as long as $E(\mathbf{U}_i) = \mathbf{0}$, so a standard linear mixed model is marginally interpretable without including an adjustment. Table 1 summarizes the form and existence of $\mathbf{d}_i^T \mathbf{a}_i$ for several common choices of link function. More specifically, this table describes the relationship between $\mathbf{x}_i^T \boldsymbol{\beta}$ and $\mathbf{d}_i^T \mathbf{a}_i$ for various link functions and random effects distributions, and also indicates whether or not there exists a closed-form solution for $\mathbf{d}_i^T \mathbf{a}_i$. The interplay between $h(\cdot)$ and f_U is explored in greater depth in Section 3.

Although we write $\mathbf{d}_i^T \mathbf{a}_i$ as a term in the conditional mean, as discussed in Section 1 we view it as part of the random effects distribution. When $\mathbf{d}_i^T \mathbf{a}_i$ depends on $\mathbf{x}_i^T \boldsymbol{\beta}$, this means that the amount each observation within a group is shifted depends on the covariates for that observation. Thus, each random effect represents a set of potential values by which to shift the observations in a group with the value for a specific observation determined by $\mathbf{x}_i^T \boldsymbol{\beta}$. This differs from the traditional formulation of a random effect and it allows one to separate systematic variation in the population, captured by $\mathbf{x}_i^T \boldsymbol{\beta}$, from individual-level variation, captured by $\mathbf{d}_i^T \mathbf{U}_i + \mathbf{d}_i^T \mathbf{a}_i$.

A situation where this new formulation of a random effect might arise is in a multilevel model where the covariates \mathbf{x}_i differ across individual units in the same group or cluster. For example, students within the same class are liable to have different characteristics. Depending on the choice of link function, the adjustment could be different for different units in the same group.

Thus, within a single group, the shift associated with the random effect for that group could vary with the measured covariates for the individual units sharing that random effect.

2.3 Dimension Reduction with Normal Random Effects

GLMMs are often implemented using normal random effects. The following proposition reduces a q -dimensional integral to a univariate one. This result is useful for computation of the adjustment in the presence of multivariate random effects, as we will demonstrate later.

Proposition 2. *For the case when $\mathbf{U}_i \sim N_q(\mathbf{0}, \Sigma)$, if the q -dimensional integral*

$$\int_{\mathbb{R}^q} h(\kappa + \mathbf{d}^T \mathbf{u} + a) \left(\frac{1}{2\pi} \right)^{\frac{q}{2}} |\Sigma|^{-1/2} \exp\left(-\frac{1}{2} \mathbf{u}^T \Sigma^{-1} \mathbf{u} \right) d\mathbf{u} \quad (7)$$

exists, then it can be expressed as a univariate integral of the form

$$\int_{\mathbb{R}} h(\kappa + v + a) \frac{1}{\sqrt{2\pi\tau^2}} \exp\left(-\frac{1}{2\tau^2} v^2 \right) dv. \quad (8)$$

3 The Form of the Adjustment

As discussed in Section 2.2 and summarized in Table 1, the form of $\mathbf{d}_i^T \mathbf{a}_i$ depends on the choice of link function and random effects distribution. This section provides more detail about the form of $\mathbf{d}_i^T \mathbf{a}_i$ for several common choices of link function.

3.1 Log Link

Consider a GLMM with a log link. That is, let the link function be $g(\cdot) = \log(\cdot)$ with the inverse link $h(\cdot) = \exp(\cdot)$. In this case, $\mathbf{d}_i^T \mathbf{a}_i$ is defined such that

$$\exp(\mathbf{x}_i^T \boldsymbol{\beta}) = \int \exp(\mathbf{x}_i^T \boldsymbol{\beta} + \mathbf{d}_i^T \mathbf{u} + \mathbf{d}_i^T \mathbf{a}_i) f_{\mathbf{U}}(\mathbf{u}) d\mathbf{u}. \quad (9)$$

Solving (9) for $\mathbf{d}_i^T \mathbf{a}_i$ leads to the following proposition:

Proposition 3. For $h(\cdot) = \exp(\cdot)$, a model of the form given by (5) and (6) is marginally interpretable if and only if $\mathbf{d}_i^T \mathbf{a}_i = -\log\{M_{\mathbf{U}}(\mathbf{d}_i)\}$, where $M_{\mathbf{U}}(\mathbf{d}_i) = E\{\exp(\mathbf{d}_i^T \mathbf{U}_i)\}$ is the moment-generating function of \mathbf{U}_i evaluated at \mathbf{d}_i .

From Proposition 3 we obtain the following corollary:

Corollary 1. For a GLMM with inverse link function $h(\cdot) = \exp(\cdot)$, an adjustment $\mathbf{d}_i^T \mathbf{a}_i$ that makes the model marginally interpretable exists if and only if $M_{\mathbf{U}}(\mathbf{d}_i)$ exists.

These results constrain the set of possible random effects distributions that can be used with this model to those with exponential tails. Consequently, the t-distribution is not a valid random effects distribution for a marginally interpretable GLMM with a log link. One could, however, use a mixture of normal distributions to approximate a t-distribution.

To better understand the role of the adjustment for a model with a log link, consider the case of a single random intercept $U_i \sim N(0, \sigma^2)$. In this case, $d_i = 1$ for all $i = 1, \dots, N$, so the adjustment $d_i a_i$ is expressed simply as a_i . From Proposition 3 we have the following result:

Corollary 2. A model of the form $E(Y_i|U_i = u) = \exp(\mathbf{x}_i^T \boldsymbol{\beta} + u + a_i)$ for which $U_i \sim N(0, \sigma^2)$ is marginally interpretable if and only if $a_i = -\sigma^2/2$ for all $i = 1, \dots, N$.

In this situation, the adjustment depends only on the random effects variance σ^2 and is independent of $\mathbf{x}_i^T \boldsymbol{\beta}$. It is simply an additive offset on the log scale that pulls the conditional mean $E(Y_i|U_i)$ down by the same amount for all $i = 1, \dots, N$. This effectively shifts the location of the random effects distribution in a manner that makes the model marginally interpretable. Since the inverse link function $\exp(\cdot)$ is convex, when the random effect has mean zero and no adjustment is made we know by Jensen's inequality that

$$E(Y_i) = E\{E(Y_i|U_i)\} = E\{\exp(\mathbf{x}_i^T \boldsymbol{\beta} + U_i)\} \geq \exp\{E(\mathbf{x}_i^T \boldsymbol{\beta} + U_i)\} = \exp(\mathbf{x}_i^T \boldsymbol{\beta}). \quad (10)$$

By pulling the conditional mean down, the adjustment counteracts the convexity of the inverse link function so that the marginal mean $E(Y_i)$ is equal to $\exp(\mathbf{x}_i^T \boldsymbol{\beta})$.

3.2 Links with Bounded Domain

Several common link functions, including the probit, logit, and complementary log-log, are defined only on a bounded subset of the real line. In turn, the range of the corresponding inverse link function $h(\cdot)$ is constrained to a bounded interval. For models with such a link function, the following theorem applies:

Theorem 1. *Consider a model of the form given in (5) with $h : \mathbb{R} \rightarrow \mathcal{I} = [\ell, u]$. Suppose $h(\cdot)$ is increasing and continuous, with $h(\zeta) \rightarrow \ell$ as $\zeta \rightarrow -\infty$ and $h(\zeta) \rightarrow u$ as $\zeta \rightarrow \infty$. Then an adjustment $\mathbf{d}_i^T \mathbf{a}_i$ that satisfies (6) exists for any choice of random effects distribution.*

Thus, one can always construct a model to be marginally interpretable when using a link function defined only on a bounded interval. We now discuss link functions with this property.

Probit Link

Let $g(\cdot) = \Phi^{-1}(\cdot)$ and $h(\cdot) = \Phi(\cdot)$, where $\Phi(\cdot)$ is the cumulative distribution function of a standard normal distribution. The range of the inverse link function $h(\cdot)$ is the bounded interval $(0, 1)$, so Theorem 1 applies and an adjustment $\mathbf{d}_i^T \mathbf{a}_i$ that makes the model marginally interpretable exists regardless of the choice of random effects distribution.

For a model with a probit link and normal random effects the adjustment has a closed form. Specifically, let $\mathbf{U}_i \sim N_q(\mathbf{0}, \Sigma)$, where Σ is a covariance matrix. One can show that without the adjustment this model would satisfy (2) with $c = (1 + \mathbf{d}_i^T \Sigma \mathbf{d}_i)^{-1/2}$ (see McCulloch et al., 2008). This leads to the following proposition:

Proposition 4. *For $h(\cdot) = \Phi(\cdot)$ and $\mathbf{U}_i \sim N_q(\mathbf{0}, \Sigma)$, a model of the form given by (5) and (6) is marginally interpretable if and only if $\mathbf{d}_i^T \mathbf{a}_i = \{(1 + \mathbf{d}_i^T \Sigma \mathbf{d}_i)^{1/2} - 1\} \mathbf{x}_i^T \boldsymbol{\beta}$.*

Thus, $\mathbf{d}_i^T \mathbf{a}_i$ is a linear function of $\mathbf{x}_i^T \boldsymbol{\beta}$. This linearity is beneficial for model interpretation, but models with a probit link are generally difficult to interpret because the probit does not have the convenient log-odds interpretation of the logit. We therefore focus on models with a logit link.

Logit Link

Consider a GLMM with link function $g(\mu) = \log\{\mu/(1 - \mu)\}$ and inverse link function $h(\eta) = \exp(\eta)/\{1 + \exp(\eta)\} = 1/\{1 + \exp(-\eta)\}$. This function $g(\cdot)$ is known both as the *logit link* and as the *logistic link*. The adjustment $\mathbf{d}_i^T \mathbf{a}_i$ for this model is defined such that

$$\frac{1}{1 + e^{-\mathbf{x}_i^T \boldsymbol{\beta}}} = \int \frac{1}{1 + e^{-(\mathbf{x}_i^T \boldsymbol{\beta} + \mathbf{d}_i^T \mathbf{u} + \mathbf{d}_i^T \mathbf{a}_i)}} f_{\mathbf{U}}(\mathbf{u}) d\mathbf{u}. \quad (11)$$

Once again, the range of the inverse link function $h(\cdot)$ is the bounded interval $(0, 1)$, so by Theorem 1 there are no restrictions on the choice of the random effects distribution. However, for most choices of random effects distribution the integral on the right-hand-side of (11) is analytically intractable and there is no closed-form solution for $\mathbf{d}_i^T \mathbf{a}_i$. One exception is the *bridge distribution* derived by Wang and Louis (2003). Provided the model contains just a single random intercept, the bridge distribution leads to a closed-form solution for $\mathbf{d}_i^T \mathbf{a}_i$ that is linear as a function of $\mathbf{x}_i^T \boldsymbol{\beta}$.

It is most common to assume that $f_{\mathbf{U}}$ is a normal density. In this article we develop a novel and efficient method for calculating $\mathbf{d}_i^T \mathbf{a}_i$ under this assumption for the random effects distribution. Our algorithm for evaluating the logistic-normal integral exploits a recursive formula developed by Pirjol (2013) that provides an exact solution to the logistic-normal integral on a specifically-defined evenly-spaced grid. Pirjol (2013) demonstrated that the integral

$$\varphi(\mu, \sigma^2) = \int \frac{1}{1 + e^w} \frac{1}{\sqrt{2\pi\sigma^2}} \exp\left\{-\frac{1}{2\sigma^2}(w - \mu)^2\right\} dw \quad (12)$$

satisfies the recursion

$$\varphi(\mu + \sigma^2, \sigma^2) = e^{-\mu - \frac{\sigma^2}{2}} \{1 - \varphi(\mu, \sigma^2)\}, \quad (13)$$

where $\varphi(0, \sigma^2) = 1/2$. See Section 4.1 for details on how this result is used to compute $\mathbf{d}_i^T \mathbf{a}_i$ accurately and efficiently.

For a model with a logit link, whether $f_{\mathbf{U}}$ is assumed to be normal or not, both the direction and magnitude of the adjustment depend on $\mathbf{x}_i^T \boldsymbol{\beta}$. In fact, the adjustment $\mathbf{d}_i^T \mathbf{a}_i$ is typically a

nonlinear function of $\mathbf{x}_i^T \boldsymbol{\beta}$. This is illustrated in Figure 1 for the case of a single normal random intercept $U_i \sim N(0, \sigma^2)$. In light of Proposition 2, the same picture would apply for q normal random effects $\mathbf{U}_i \sim N_q(\mathbf{0}, \boldsymbol{\Sigma})$ if we were to replace σ^2 with $\mathbf{d}_i^T \boldsymbol{\Sigma} \mathbf{d}_i$. The direction of the adjustment is driven by the convexity of the inverse link function. The function $h(\eta)$ is convex for $\eta < 0$ and concave for $\eta > 0$, so the adjustment is negative when $\mathbf{x}_i^T \boldsymbol{\beta} < 0$ and positive when $\mathbf{x}_i^T \boldsymbol{\beta} > 0$. It is also evident from Figure 1 that the magnitude of $\mathbf{d}_i^T \mathbf{a}_i$ is increasing in both σ^2 and $|\mathbf{x}_i^T \boldsymbol{\beta}|$. For very large $\mathbf{x}_i^T \boldsymbol{\beta}$ we have the following result:

Proposition 5. *For $h(\cdot) = \text{logit}^{-1}(\cdot)$ and $\mathbf{U}_i \sim N_q(\mathbf{0}, \boldsymbol{\Sigma})$, the value of $\mathbf{d}_i^T \mathbf{a}_i$ that allows a model of the form given by (5) to satisfy (6) converges to $\frac{1}{2} \mathbf{d}_i^T \boldsymbol{\Sigma} \mathbf{d}_i \times \text{sign}(\mathbf{x}_i^T \boldsymbol{\beta})$ as $|\mathbf{x}_i^T \boldsymbol{\beta}| \rightarrow \infty$.*

Figure 1 also shows that, for a model with a logit link, observations that have different values of the covariates \mathbf{x}_i also have different adjustments. This helps to illustrate the point from Section 2.2 that with the logit link, observations on individual units that share a random effect but have different measured covariates are not shifted by the same amount. Rather, the magnitude of the shift associated with the random effect for each observation is determined by the value of $\mathbf{x}_i^T \boldsymbol{\beta}$ for that observation.

Complementary Log-Log Link

Consider a GLMM with $g(\mu) = \log\{-\log(1 - \mu)\}$ and $h(\eta) = 1 - \exp\{-\exp(\eta)\}$. The adjustment $\mathbf{d}_i^T \mathbf{a}_i$ for this model is defined such that

$$\exp\{-\exp(\mathbf{x}_i^T \boldsymbol{\beta})\} = \int \exp\{-\exp(\mathbf{x}_i^T \boldsymbol{\beta} + \mathbf{d}_i^T \mathbf{u} + \mathbf{d}_i^T \mathbf{a}_i)\} f_{\mathbf{U}}(\mathbf{u}) d\mathbf{u}. \quad (14)$$

As with a logit link, there are no restrictions on the choice of the random effects distribution, but in most cases there is no closed-form solution for $\mathbf{d}_i^T \mathbf{a}_i$. Wang and Louis (2003) also derived a bridge distribution for this link function that leads to a closed-form adjustment that is linear in $\mathbf{x}_i^T \boldsymbol{\beta}$. For more conventional choices of random effects distribution, namely normal random effects, one must use some form of numerical integration, such as Gaussian quadrature, to calculate $\mathbf{d}_i^T \mathbf{a}_i$.




Figure 1: Plot of the adjustment a as a function of κ for various values of σ , where $\kappa = \mathbf{x}_i^T \boldsymbol{\beta}$ represents the fixed effects portion of the model.

3.3 Links with Range Restrictions

A number of common link functions map into a proper subset of the real line and therefore require conditions on $\mathbf{x}_i^T \boldsymbol{\beta}$ to ensure that the model is defined. For example, the square root transformation is typically defined to have nonnegative range and no real number has a reciprocal of zero. Additive random effects with support on the entire real line could lead to problems in models with these link functions because $\mathbf{x}_i^T \boldsymbol{\beta} + \mathbf{d}_i^T \mathbf{U}_i$ could fall outside the domain of the inverse link function $h(\cdot)$. Thus, special care must be taken with these link functions, as described below.

Square Root Link

Consider a GLMM with link function $g(\mu) = \mu^{1/2}$ and inverse link function $h(\eta) = \eta^2$. For such a model one typically includes the restriction that $\mathbf{x}_i^T \boldsymbol{\beta} \geq 0$. Including the adjustment, we adopt the restriction that $\mathbf{x}_i^T \boldsymbol{\beta} + \mathbf{d}_i^T \mathbf{a}_i \geq 0$. The adjustment is defined such that

$$(\mathbf{x}_i^T \boldsymbol{\beta})^2 = \int (\mathbf{x}_i^T \boldsymbol{\beta} + \mathbf{d}_i^T \mathbf{u} + \mathbf{d}_i^T \mathbf{a}_i)^2 f_{\mathbf{U}}(\mathbf{u}) d\mathbf{u}. \quad (15)$$

If we assume $E(\mathbf{U}_i) = \mathbf{0}$, then (15) reduces to

$$(\mathbf{x}_i^T \boldsymbol{\beta})^2 = (\mathbf{x}_i^T \boldsymbol{\beta} + \mathbf{d}_i^T \mathbf{a}_i)^2 + \text{Var}(\mathbf{d}_i^T \mathbf{U}_i), \quad (16)$$

which is quadratic in $\mathbf{d}_i^T \mathbf{a}_i$ and leads to the following result:

Proposition 6. *For $h(\eta) = \eta^2$ and $E(\mathbf{U}_i) = \mathbf{0}$, a model of the form given by (5) and (6) subject to the restriction that $\mathbf{x}_i^T \boldsymbol{\beta} + \mathbf{d}_i^T \mathbf{a}_i \geq 0$ is a marginally interpretable model if and only if $\mathbf{d}_i^T \mathbf{a}_i = -\mathbf{x}_i^T \boldsymbol{\beta} + \{(\mathbf{x}_i^T \boldsymbol{\beta})^2 - \text{Var}(\mathbf{d}_i^T \mathbf{U}_i)\}^{1/2}$.*

Thus, for a model with a square root link, $\mathbf{d}_i^T \mathbf{a}_i$ is a nonlinear function of $\mathbf{x}_i^T \boldsymbol{\beta}$ and is only defined when $\mathbf{x}_i^T \boldsymbol{\beta} \geq \{\text{Var}(\mathbf{d}_i^T \mathbf{U}_i)\}^{1/2}$. If the random effects variance is too large, the model cannot be fit.

Reciprocal Link

Finally, consider a model for which $g(\mu) = h(\mu) = 1/\mu$. For a fixed effects model with this link one typically includes the restriction that $\mathbf{x}_i^T \boldsymbol{\beta} > 0$. When a random intercept U_i is included in the model, the fact that $h(\cdot)$ tends to infinity as its argument approaches zero forces us to also include restrictions on the distribution of U_i . In particular, we want a model for which

$$\frac{1}{\mathbf{x}_i^T \boldsymbol{\beta}} = \int \frac{1}{\mathbf{x}_i^T \boldsymbol{\beta} + u} f_U(u) du, \quad (17)$$

so f_U must be defined such that the integral on the right-hand-side of (17) exists. This restriction forces us to move away from models of the form given in (5). Rather than adjusting the location of the random effect based on each individual's observed covariates, we must alter the shape of the distribution of the random effect based on the observed covariates.

One distribution for U_i that allows us to satisfy (17) is a shifted gamma distribution. Specifically, let $\mathbf{x}_i^T \boldsymbol{\beta} + U_i$ follow a gamma distribution with shape parameter α_i and rate parameter β_i so that $E(U_i) = \alpha_i \beta_i - \mathbf{x}_i^T \boldsymbol{\beta}$. Then the integral on the right-hand-side of (17) is equal to $\{\beta_i(\alpha_i - 1)\}^{-1}$. Thus, $\mathbf{x}_i^T \boldsymbol{\beta} = \beta_i(\alpha_i - 1)$, so by placing additional conditions on α_i and β_i one can determine the appropriate gamma distribution for U_i for each $\mathbf{x}_i^T \boldsymbol{\beta}$. Alternatively, one could let $\mathbf{x}_i^T \boldsymbol{\beta} + U_i$ follow

an inverse gamma distribution with parameters α_i and β_i , and be constrained by the relationship $\mathbf{x}_i^T \boldsymbol{\beta} = (\alpha_i \beta_i)^{-1}$. In either case it is the shape, not the location, of the random effects distribution that varies with $\mathbf{x}_i^T \boldsymbol{\beta}$ in this marginally interpretable model.

4 Computation

Since the adjustment that we propose is a deterministic function of $\mathbf{x}_i^T \boldsymbol{\beta}$ and the parameters characterizing $f_{\mathbf{U}}$ (which we denote $\boldsymbol{\alpha}$), it can easily be incorporated into techniques that are commonly used to fit conventional GLMMs. Taking a Bayesian approach, for example, one might employ a Metropolis-Hastings MCMC algorithm to generate samples from the posterior distribution of the unknown parameters. Such an algorithm involves iteratively proposing values for the unknown quantities $\boldsymbol{\beta}$, $\boldsymbol{\alpha}$, and \mathbf{U} and choosing to either accept or reject those values in an effort to sample from the posterior density $\pi(\boldsymbol{\beta}, \boldsymbol{\alpha}, \mathbf{U} | \mathbf{Y})$. Denoting $\boldsymbol{\theta} = (\boldsymbol{\beta}, \boldsymbol{\alpha}, \mathbf{U})^T$, the acceptance probability at the t^{th} iteration is given by

$$\min \left\{ 1, \frac{f_{\mathbf{Y}|\boldsymbol{\theta}}(\mathbf{Y}|\boldsymbol{\theta}^*)\pi_{\boldsymbol{\theta}}(\boldsymbol{\theta}^*)q(\boldsymbol{\theta}^{(t)}|\boldsymbol{\theta}^*)}{f_{\mathbf{Y}|\boldsymbol{\theta}}(\mathbf{Y}|\boldsymbol{\theta}^{(t)})\pi_{\boldsymbol{\theta}}(\boldsymbol{\theta}^{(t)})q(\boldsymbol{\theta}^*|\boldsymbol{\theta}^{(t)})} \right\}, \quad (18)$$

where $\boldsymbol{\theta}^{(t)}$ is the current state of $\boldsymbol{\theta}$, $\boldsymbol{\theta}^*$ is the proposed state of $\boldsymbol{\theta}$, $f_{\mathbf{Y}|\boldsymbol{\theta}}(\cdot)$ is the conditional density of \mathbf{Y} , $\pi_{\boldsymbol{\theta}}(\cdot)$ is the prior density for $\boldsymbol{\theta}$, and $q(\cdot)$ is the proposal density. Within each update of $\boldsymbol{\theta}$ the proposals $\boldsymbol{\beta}^*$ and $\boldsymbol{\alpha}^*$ can be used along with \mathbf{x}_i and \mathbf{d}_i , which are treated as fixed and known, to compute the adjustment $\mathbf{d}_i^T \mathbf{a}_i$, which can in turn be included in evaluation of the likelihood.

The computational expense added to the algorithm by including the adjustment is driven by how difficult it is to compute $\mathbf{d}_i^T \mathbf{a}_i$, which varies based on the choice of link function and random effects distribution. When $\mathbf{d}_i^T \mathbf{a}_i$ has a closed form, inclusion of the adjustment has little impact on the amount of computation required, whereas in situations lacking a closed-form solution there could be considerable computational expense associated with calculating $\mathbf{d}_i^T \mathbf{a}_i$. The most common situation without a closed-form solution is a GLMM with a logit link and normal random effects. An efficient approach for computing the adjustment in such a case is described below.

4.1 Efficient and Accurate Evaluation of the Logistic-Normal Integral

Although the logistic-normal integral is analytically intractable, several numerical approaches exist for evaluating it. Common approaches include Gaussian quadrature and adaptive quadrature schemes. There also exist a number of algorithms tailored specifically to evaluating the logistic-normal integral, including methods proposed by Crouch and Spiegelman (1990) and Monahan and Stefanski (1992). The method of Monahan and Stefanski (1992) involves approximating the inverse logit function $h(z)$ with a weighted mixture of normal distributions

$$h_k^*(z) = \sum_{i=1}^k p_{k,i} \Phi(z s_{k,i}), \quad (19)$$

where $\Phi(\cdot)$ is the cumulative distribution function of a standard normal distribution and the weights $p_{k,i}$ and $s_{k,i}$ are chosen to minimize the maximum approximation error over all values of z . This leads to the integral approximation

$$\int h(z) \frac{1}{\sigma} \phi\left(\frac{z - \xi}{\sigma}\right) dz \approx \int h_k^*(z) \frac{1}{\sigma} \phi\left(\frac{z - \xi}{\sigma}\right) dz = \sum_{i=1}^k p_{k,i} \Phi\left\{ \frac{\xi s_{k,i}}{(1 + \sigma^2 s_{k,i}^2)^{1/2}} \right\}, \quad (20)$$

which is within 2.1×10^{-9} of the true value of the integral for all values of κ and σ when one uses $k = 8$ mixture weights. One could use fewer than eight mixture weights to improve computational efficiency, but the increase in speed from using fewer weights is small relative to the corresponding loss of accuracy, so we recommend using $k = 8$.

More recently, Pirjol (2013) developed the recursive formula given in (13), which provides an exact solution to the logistic-normal integral on a specific evenly-spaced grid. To simplify notation we denote $\kappa = \mathbf{x}_i^T \boldsymbol{\beta}$ and $a = a_i$, and use $\varphi(\cdot, \cdot)$ as defined in (12). Given κ and σ^2 , the equation we must solve for a when $g(\cdot) = \text{logit}(\cdot)$ and $U_i \sim N(0, \sigma^2)$ is

$$h(\kappa) = 1 - \varphi(\kappa + a, \sigma^2). \quad (21)$$

Without loss of generality, due to the symmetry of the problem, we need only consider the case of $\kappa > 0$. When $\kappa < 0$ the adjustment has the same magnitude but opposite sign as if $\kappa = |\kappa|$. Further, as a consequence of Proposition 2, the strategy described here for the univariate case also applies when one has multivariate normal random effects.

Several different techniques, such as binary segmentation or a Newton-Raphson algorithm, can be used to solve (21) for a . Any such technique requires evaluating $\varphi(\kappa + a^*, \sigma^2)$ for several potential values a^* of the adjustment a . While we can use (13) to calculate $\varphi(t\sigma^2, \sigma^2)$ exactly for any integer t , it is unlikely that the desired a will be such that $\kappa + a$ is an integer multiple of σ^2 . Thus, we have need of an approximate numerical integration procedure. Since the function $\varphi(\cdot, \cdot)$ is decreasing in its first argument, we can use (13) to quickly identify an interval of length σ^2 in which $\kappa + a$ must reside. We denote this interval $(t^*\sigma^2, (t^* + 1)\sigma^2)$, where t^* is a nonnegative integer. By narrowing our search for the correct value of a to such an interval we cut down the required number of evaluations of $\varphi(\kappa + a^*, \sigma^2)$.

We also use the recursive result of Pirjol (2013) to improve the accuracy of the necessary integral approximations. After identifying t^* , we employ binary segmentation to search within the interval $(t^*\sigma^2, (t^* + 1)\sigma^2)$ for the value of $\kappa + a$ satisfying (21). To evaluate $\varphi(\kappa + a^*, \sigma^2)$ for $\kappa + a^* \in (t^*\sigma^2, (t^* + 1)\sigma^2)$ we could simply use (20). However, we find that the Monahan-Stefanski approximation is generally more accurate near zero than away from it. We therefore use (20) to compute $\varphi(\kappa + a^* - t^*\sigma^2, \sigma^2)$ and then apply (13) t^* times to obtain $\varphi(\kappa + a^*, \sigma^2)$.

Our method for evaluating the logistic-normal integral should ideally provide more accurate approximations than competing methods that are comparable in terms of computational efficiency. To assess the speed and accuracy of our approach we compared it to both 30-point Gaussian quadrature and to a direct application of (20). Specifically, for each of the 80 values of σ in the set $\{0.05, 0.10, \dots, 4.00\}$ we evaluated the integral $1 - \varphi(\mu, \sigma^2)$ for 1,000 values of μ in each of the four intervals $[0, \sigma^2]$, $[\sigma^2, 2\sigma^2]$, $[2\sigma^2, 3\sigma^2]$, and $[3\sigma^2, 4\sigma^2]$ using our method, the method of Monahan and Stefanski, 30-point Gauss-Hermite quadrature, and 1,000-point Gauss-Hermite quadrature.




Figure 2: Maximum error relative to 1,000-point Gauss-Hermite quadrature for various integral approximations in the intervals $[\sigma^2, 2\sigma^2]$ (left panel) and $[2\sigma^2, 3\sigma^2]$ (right panel). Machine accuracy is approximately 10^{-16} , accounting for the floor in the two plots.

ture. This required 4,000 integral evaluations for each of the 80 values of σ and each method. These evaluations were completed on a Dual Quad Core Xeon 2.66 E5430 computer with 32 gigabytes of RAM. To ensure a fair comparison of speed, all four approaches were implemented using the Rcpp package in R (R Core Team, 2016; Eddelbuettel and François, 2011; Eddelbuettel, 2013). Gaussian quadrature with 1,000 quadrature points was treated as the *gold standard* to which the other three methods were compared to assess accuracy.

For each of the competing methods and each value of σ we computed the maximum “error” relative to 1,000-point quadrature within each of the four intervals for μ . Figure 2 summarizes the results of the accuracy assessment for $[\sigma^2, 2\sigma^2]$ and $[2\sigma^2, 3\sigma^2]$. Although 30-point Gaussian quadrature is the most accurate for small values of σ , our approach is the most accurate in the majority of cases. Notably, our approach, which combines the recursion of Pirjol (2013) with the Monahan-Stefanski approximation, is uniformly no worse than a direct application of (20).

The 320,000 integral evaluations required for the accuracy assessment took 2.1 seconds for our approach compared to 2.1 seconds for the direct application of the Monahan-Stefanski method, 2.2 seconds for 30-point Gaussian quadrature, and 19.4 seconds for 1,000-point Gaussian quadrature. Thus, the efficiency of our approach is comparable to that of the Monahan-Stefanski approach

and slightly better than that of 30-point Gaussian quadrature. Further, 1,000-point quadrature is considerably less efficient than the other three methods. We conclude that our method offers the best tradeoff between accuracy and efficiency.

4.2 Improving Mixing in the Presence of Many Random Effects

A challenge associated with using MCMC to sample from a high-dimensional posterior density is poor mixing. Due to the large number of unknown parameters, the only proposals that get accepted are those representing relatively small steps from the current state of the Markov chain. Consequently, there is substantial autocorrelation in the Markov chain and it is necessary to run the algorithm for an exceedingly long time to generate samples that represent approximately independent draws from the target posterior.

One way to improve mixing is to sample the parameters in blocks, but this may not always be enough. We have found that when there are many random effects the latent random variables can dominate the likelihood and cause very few proposed β^* to be accepted. To overcome this, we suggest that with each proposed β^* one simultaneously proposes random effects U^* that are consistent with the proposed fixed effects in order to increase the chance of acceptance. An example illustrating this approach is given in Section 5.2. The β^* and U^* are proposed in such a manner that there is no net impact on the likelihood and the decision to accept or reject the proposed values is based entirely on the prior distributions. A formal update of U is still required, but this strategy improves the acceptance rate for β and thereby facilitates faster mixing.

5 Applications

In this section, we provide two examples of marginally interpretable GLMMs applied to the analysis of real data. First, we use the technique described in Section 4.1 to fit a model with a logit link to data from a teratological experiment on rats. Next, we use the strategy introduced in

Section 4.2 to sample from the posterior of a marginally interpretable GLMM with a log link for data from a clinical trial of epileptics.

5.1 Rat Teratology

Our first example comes from a teratological experiment on rats conducted by Weil (1970). A group of 16 female rats was fed a diet containing a chemical agent during pregnancy and lactation, while another group of 16 female rats was fed a control diet. Counts were made of the number of pups in each litter to survive four days from birth and to survive the 21-day lactation period. Interest lies in the proportion of pups to survive 21 days among those alive after four days.

For $i = 1, 2$ and $j = 1, \dots, 16$, we denote the number of pups in litter j receiving treatment i to survive four days by m_{ij} , the number of pups to survive 21 days by Y_{ij} , and the proportion of pups to survive 21 days by $p_{ij} = Y_{ij}/m_{ij}$. We include in our model a fixed effect for the treatment ($x_1 = 1$ for the treatment group, $x_2 = -1$ for the control group) and random effects for litter (denoted U_{ij}). Earlier analyses of these data established that there is more between-litter heterogeneity in the treatment group than in the control group (see Liang and Hanfelt, 1994; Heagerty and Zeger, 2000; Wang and Louis, 2004). We therefore allow different random effects variances for the two treatment groups. We assume $U_{ij} \stackrel{\text{ind}}{\sim} N(0, \sigma_i^2)$ and $Y_{ij}|\boldsymbol{\beta}, U_{ij} \stackrel{\text{ind}}{\sim} \text{Binomial}\{m_{ij}, E(p_{ij}|\boldsymbol{\beta}, U_{ij})\}$, and model the conditional mean as

$$E(p_{ij}|\boldsymbol{\beta}, U_{ij}) = h(\beta_0 + \beta_1 x_i + U_{ij} + a_i), \quad (22)$$

where $h(\cdot)$ is the inverse logit function, $\boldsymbol{\beta} = (\beta_0, \beta_1)^T$ is the vector of fixed effects parameters, and a_i is the adjustment that ensures the model is marginally interpretable. Since all litters in the same treatment group have the same covariate x_i and the same random effects variance σ_i^2 , they also have the same adjustment a_i .

We adopt a Bayesian approach and use MCMC to sample from the posterior distribution of the unknown parameters in our model. Our prior distributions for β_0 , β_1 , $\log(\sigma_1^2)$, and $\log(\sigma_2^2)$ are $N(0, 25)$, $N(0, 10)$, $N(-1/2, 1)$, and $N(-1/2, 1)$, respectively. To sample from our target posterior

Table 2: Posterior means of the unknown parameters in the model for the rat teratology data (with corresponding posterior standard deviations in parentheses)

Parameter	Marginally Interpretable Model	Conventional GLMM
β_0	1.66 (0.24)	1.99 (0.31)
β_1	-0.51 (0.23)	-0.39 (0.31)
σ_1	1.54 (0.41)	1.60 (0.43)
σ_2	0.73 (0.29)	0.75 (0.30)

we iteratively update blocks of parameters using Metropolis steps. We first update $\beta = (\beta_0, \beta_1)^T$, then $\alpha = (\sigma_1^2, \sigma_2^2)^T$, and finally $\mathbf{U} = (U_{1,1}, \dots, U_{1,16}, U_{2,1}, \dots, U_{2,16})^T$. When necessary, we compute the adjustment a_i using the technique described in Section 4.1. This MCMC algorithm was carried out both with the adjustment included in the model and without it. Each chain was run for 1,010,000 steps, with the first 10,000 steps discarded as burn-in and every 100th step thereafter retained for the final sample. This resulted in 10,000 draws from the posterior distribution for each model. Additional details regarding the algorithm are provided in the Supplementary Material.

Table 2 provides posterior means and standard deviations for the parameters in both the marginally interpretable model and the conventional GLMM, which does not include the adjustment. Figure 3 displays kernel density estimates based on the posterior samples for the two models. Notably, for β_1 , which corresponds to the treatment effect, the tail area above zero is 0.016 for the marginally interpretable model. This is considerably less than the tail area of 0.101 for the conventional GLMM. Thus, many would draw different conclusions about the importance of the treatment effect using the two different models. Indeed, Bayes factors for a test of no treatment effect ($H_0 : \beta_1 = 0$), computed using the Savage-Dickey density ratio (see Dickey, 1971; Verdinelli and Wasserman, 1995), favor the null hypothesis and come in at 1.27 for the marginally interpretable model and 4.41 for the conventional GLMM, thereby confirming this disparity.

Since more rat pups survive than do not, the log-odds of survival are generally positive and therefore fall in a region where the inverse link function $h(\cdot)$ is concave. Consequently, in the presence of random effects, the marginal mean is pulled down relative to the conditional mean.




Figure 3: Kernel density estimates of the posterior densities for the unknown parameters in the model for the rat teratology data. Estimates obtained from the marginally interpretable model are in black while those obtained from the conventional GLMM are in gray.

This explains why the marginally interpretable intercept is less than β_0 in the conventional GLMM. Further, since the random effects variance is greater in the treatment group than in the control group, the concavity of $h(\cdot)$ has a greater impact on the treatment group. This contributes to β_1 , the coefficient for the treatment effect, having a lesser value in the marginally interpretable model than in the conventional GLMM. Note that the marginally interpretable β_1 is not attenuated toward zero relative to β_1 from the conventional GLMM, as one might expect, because the random intercept is not independent of x_i .

We compare the expected 21-day survival rates between the two treatment groups. For the marginally interpretable model, the expected proportion of rat pups in the treatment group to survive 21 days among those alive after four days is $E(p|\beta, \alpha, x = 1) = h(\beta_0 + \beta_1)$ and in the control group is $E(p|\beta, \alpha, x = -1) = h(\beta_0 - \beta_1)$. For the conventional GLMM the same expectation is $E(p|\beta, \alpha, x = 1) = \int h(\beta_0 + \beta_1 + u)f_U(u)du$ for the treatment group and $E(p|\beta, \alpha, x = -1) = \int h(\beta_0 - \beta_1 + u)f_U(u)du$ for the control group. Note that the parameters α enter this expression through the random effects distribution f_U . Interest lies in whether or




Figure 4: Kernel density estimates of the posterior density of the difference in the expected 21-day survival rate between the two treatment groups based on the marginally interpretable model (black) and the conventional GLMM (gray).

not the quantity $E(p|\beta, \alpha, x = 1) - E(p|\beta, \alpha, x = -1)$ is nonzero. Kernel density estimates of the posterior density for this quantity under the two models are shown in Figure 4. The integral evaluation required for the conventional GLMM was accomplished using Monte Carlo integration. Under both models, most of the posterior mass is below zero. For the marginally interpretable model, the tail area above zero is 0.016, matching the tail area for β_1 . However, for the conventional GLMM, the tail area above zero is 0.041, which contrasts sharply with the tail area of 0.101 for β_1 .

5.2 Epileptic Seizures

Our second example comes from a clinical trial of 59 epileptics conducted by Leppik et al. (1987). Each subject received either a placebo or the drug progabide and then made four successive follow-up visits to the clinic during which they reported the number of partial seizures they had suffered in the two-week period immediately preceding the visit. We denote these reported counts by Y_{ij} , where $i = 1, \dots, 59$ indexes the subjects and $j = 1, 2, 3, 4$ indexes the visits. Thall and Vail (1990) used GEE to fit a marginal model to these data. They included as predictors the logarithm of one-fourth of the baseline count of partial seizures suffered by each patient in the eight-week period prior to treatment (denoted $BASE_i$), a treatment indicator (1 if progabide, 0 if placebo, denoted TRT_i), the interaction between $BASE_i$ and TRT_i , the logarithm of the subject's age in

years (denoted AGE_i), and a fourth-visit indicator (1 for the subject's fourth post-treatment visit, 0 otherwise, denoted VISIT4_j). Breslow and Clayton (1993) and Gamerman (1997) fit a GLMM with the same fixed effects and also two levels of random effects. Their model has the form

$$\begin{aligned} E(Y_{ij}|\boldsymbol{\beta}, \gamma_i, \delta_{ij}) = & \exp(\beta_0 + \beta_1 \times (\text{BASE}_i) + \beta_2 \times (\text{TRT}_i) + \beta_3 \times (\text{BASE}_i * \text{TRT}_i) + \\ & \beta_4 \times (\text{AGE}_i) + \beta_5 \times (\text{VISIT4}_j) + \gamma_i + \delta_{ij}), \end{aligned} \quad (23)$$

where $\boldsymbol{\beta} = (\beta_0, \beta_1, \dots, \beta_5)^T$ is the vector of fixed effects parameters, the $\gamma_i \stackrel{\text{ind}}{\sim} N(0, \sigma^2)$ are random subject effects, and the $\delta_{ij} \stackrel{\text{ind}}{\sim} N(0, \tau^2)$ are random effects for visit within subject. Further, conditional on the random effects γ_i and δ_{ij} , the reported seizure counts are assumed to be independent observations from a $\text{Poisson}\{E(Y_{ij}|\boldsymbol{\beta}, \gamma_i, \delta_{ij})\}$. Breslow and Clayton (1993) fit this model using PQL while Gamerman (1997) used MCMC to sample from a Bayesian version of this model.

We adopt a Bayesian approach and sample from a mixed model analogous to (23), but include an adjustment to ensure that the model is marginally interpretable. In light of Section 3.1, the adjustment is simply $a_{ij} = -\sigma^2/2 - \tau^2/2$ for all i and j . We place $N_6(\mathbf{0}, 100\mathbf{I}_6)$, $N(-1, 2)$, and $N(-1, 2)$ prior distributions on $\boldsymbol{\beta}$, $\log(\sigma^2)$, and $\log(\tau^2)$, respectively. To sample from our target posterior using MCMC, the vector of parameters we must update is $\boldsymbol{\theta} = (\boldsymbol{\beta}, \boldsymbol{\alpha}, \boldsymbol{\gamma}, \boldsymbol{\delta})^T$, where $\boldsymbol{\beta}$ is the vector of fixed effects parameters, $\boldsymbol{\alpha} = (\sigma^2, \tau^2)^T$ includes the parameters characterizing the random effects distribution, $\boldsymbol{\gamma} = (\gamma_1, \dots, \gamma_{59})^T$ includes the 59 latent variables associated with the subject random effect, and $\boldsymbol{\delta} = (\delta_{1,1}, \dots, \delta_{59,4})^T$ includes the 236 latent variables associated with the visit random effect. Due to the presence of 295 latent variables in this model, proposals for $\boldsymbol{\beta}$ are rarely accepted when we use a basic MCMC algorithm that employs Metropolis steps to update the parameters in blocks.

To address the problem with slow mixing, we simultaneously propose $\boldsymbol{\gamma}^*$ and $\boldsymbol{\delta}^*$ to be consistent with each proposed $\boldsymbol{\beta}^*$ as described in Section 4.2. Specifically, for each $\boldsymbol{\beta}^*$ we also propose the following γ_i^* and δ_{ij}^* for each $i = 1, \dots, 59$ and $j = 1, 2, 3, 4$:

$$\gamma_i^* = \gamma_i^{(t)} + \mathbf{x}_{i,1}^T (\boldsymbol{\beta}^{(t)} - \boldsymbol{\beta}^*), \quad (24)$$

$$\delta_{i,4}^* = \delta_{i,4}^{(t)} + (\mathbf{x}_{i,4}^T - \mathbf{x}_{i,1}^T)(\boldsymbol{\beta}^{(t)} - \boldsymbol{\beta}^*), \text{ and } \delta_{ij}^* = \delta_{ij}^{(t)} \text{ for } j = 1, 2, 3. \quad (25)$$

We then choose to accept or reject $\boldsymbol{\beta}^*$, $\boldsymbol{\gamma}^*$, and $\boldsymbol{\delta}^*$ collectively and set $\boldsymbol{\beta}^{(t+1)}$, $\boldsymbol{\gamma}'$, and $\boldsymbol{\delta}'$ accordingly. The intermediate states $\boldsymbol{\gamma}'$ and $\boldsymbol{\delta}'$ are used in place of $\boldsymbol{\gamma}^{(t)}$ and $\boldsymbol{\delta}^{(t)}$ until $\boldsymbol{\gamma}$ and $\boldsymbol{\delta}$ are formally updated. Proposing random effects to be consistent with the fixed effects in this manner increases the acceptance rate for $\boldsymbol{\beta}$ from 13.0% to 51.2%. Using this improved proposal scheme also decreases the integrated autocorrelation time for β_3 , which corresponds to the interaction effect and has the highest such value among the six fixed effects parameters, from 580.0 to 165.6.

In summary, to sample from the posterior distribution of the unknown parameters in our model we repeatedly perform the following updates:

1. Update $\boldsymbol{\alpha}^{(t+1)}$ given $\boldsymbol{\beta}^{(t)}$, $\boldsymbol{\gamma}^{(t)}$, and $\boldsymbol{\delta}^{(t)}$;
2. Update $\boldsymbol{\beta}^{(t+1)}$ given $\boldsymbol{\alpha}^{(t+1)}$, $\boldsymbol{\gamma}^{(t)}$, $\boldsymbol{\delta}^{(t)}$, while also setting $\boldsymbol{\gamma}'$ and $\boldsymbol{\delta}'$;
3. For $i = 1, \dots, 59$, update $\gamma_i^{(t+1)}$ given $\boldsymbol{\beta}^{(t+1)}$, $\boldsymbol{\alpha}^{(t+1)}$, $\boldsymbol{\gamma}'_{-i}$, and $\boldsymbol{\delta}'$, where $\boldsymbol{\gamma}'_{-i}$ represents all elements of $\boldsymbol{\gamma}$ except γ_i ;
4. For $i = 1, \dots, 59$ and $j = 1, 2, 3, 4$, update $\delta_{ij}^{(t+1)}$ given $\boldsymbol{\beta}^{(t+1)}$, $\boldsymbol{\alpha}^{(t+1)}$, $\gamma_i^{(t+1)}$, and $\boldsymbol{\delta}'_{-(ij)}$, where $\boldsymbol{\delta}'_{-(ij)}$ represents all elements of $\boldsymbol{\delta}$ except δ_{ij} .

We carried out this MCMC algorithm both for the marginally interpretable model and the conventional GLMM. We ran each chain for 2,100,000 steps, discarding the first 100,000 steps as burn-in and retaining every 200th step thereafter to obtain a final sample of 10,000 draws from the posterior distribution for each model. Additional details are provided in the Supplementary Material.

Table 3 displays parameter estimates with corresponding measures of uncertainty for the marginally interpretable model and the conventional GLMM. With the exception of the intercept β_0 , the two sets of parameter estimates are virtually identical. Breslow and Clayton (1993) noted that the slope parameters in this model have both a marginal and conditional interpretation while Ritz and Spiegelman (2004) stated that this will generally be the case for a model with a log link

Table 3: Posterior means of the unknown parameters in the model for the epilepsy data (with corresponding posterior standard deviations in parentheses)

Parameter	Marginally Interpretable Model	Conventional GLMM
β_0	-1.19 (1.23)	-1.38 (1.23)
β_1	0.88 (0.14)	0.88 (0.14)
β_2	-0.95 (0.42)	-0.96 (0.43)
β_3	0.35 (0.22)	0.35 (0.22)
β_4	0.48 (0.36)	0.48 (0.36)
β_5	-0.10 (0.09)	-0.10 (0.09)
σ	0.50 (0.07)	0.50 (0.07)
τ	0.37 (0.04)	0.37 (0.04)

and a random intercept that is independent of the covariates in the model. The intercept for the marginally interpretable model is greater than the intercept for the conventional GLMM due to the tendency of the convex inverse link function to pull the marginal mean up.

What separates the model with the adjustment from the conventional GLMM is its marginal interpretation. Suppose we were interested in the average expected seizure count across all subjects in the population with a particular set of covariates. For the marginally interpretable model, the quantity of interest is simply $E(Y|\beta, \alpha) = \exp(\mathbf{x}^T \beta)$. Note that this quantity does not functionally depend on the parameters α that characterize the random effects distribution. For the conventional GLMM, the marginal mean is $E(Y|\beta, \alpha) = \int \exp(\mathbf{x}^T \beta + \mathbf{d}^T \mathbf{u}) f_{\mathbf{U}}(\mathbf{u}) d\mathbf{u}$, which does depend on the random effects distribution. Fixed effects should be stable across different samples from the same population and perturbations of the random effects distribution should not impact them. If one were to collect a different sample of epileptics, the estimates for σ^2 and τ^2 would likely change, but the fixed effects β should remain unchanged. Estimates of the fixed effects in a conventional GLMM will thus be more sensitive to the random effects than analogous estimates in a marginally interpretable model. The marginally interpretable model is therefore more generalizable to the entire population of interest than the conventional GLMM.

The marginal model that Thall and Vail (1990) fit via GEE can also be used to make marginal inferences. However, obtaining subject-specific predictions is considerably more difficult with

such a model. The marginally interpretable GLMM is the only model from which one can easily obtain both individual-level predictions and generalizable estimates of the marginal mean. It provides a single, unified model that can be interpreted either marginally or conditionally depending on the goals of one’s analysis.

Finally, we note that the model we have fit here could potentially be simplified. There is evidence that the fourth-visit indicator, age, and the interaction between baseline count and treatment could all be removed from the model. We chose to leave these potentially unnecessary predictors in the model to remain consistent with earlier analyses of the same data.

6 Conclusion and Discussion

In this article we have defined a class of marginally interpretable GLMMs and described the form of the adjustment that appears in these models for several commonly-used link functions. Unlike conventional GLMMs, which must be interpreted conditional on the random effects, these marginally interpretable GLMMs preserve the marginal mean even when the link function is nonlinear. Consequently, model parameters can be given a *population-averaged* interpretation as well as a *subject-specific* or *cluster-specific* interpretation. In this sense, marginally interpretable GLMMs are comparable to marginal models fit via GEE, but unlike a purely marginal model, a marginally interpretable GLMM is a fully-specified model with a density for the data that can be used to make individual-level predictions in addition to marginal inferences. We have also provided details regarding how to fit marginally interpretable GLMMs, including a fast and accurate algorithm for computing the logistic-normal integral.

Many of the examples we have provided relate to models with normal random effects, in part because the normal distribution is a common choice for random effects distributions. However, a marginally interpretable GLMM does not require normal random effects and the methods in this paper apply to a wide array of random effects distributions. One interesting class of random effects distributions consists of mixtures of normal distributions. Mixed models that represent the random

effects distribution as a mixture of normals (see Magder and Zeger, 1996; Caffo et al., 2007) allow considerable flexibility in the shape of the random effects distribution.

In this article we have focused on deriving the adjustment in a marginally interpretable GLMM by relating the conditional mean to the marginal mean. Acknowledging that the marginal mean may not always be of interest, we can define marginally interpretable models in other mixed model settings. For example, in mixed-effect quantile regression models (see Koenker, 2004; Geraci and Bottai, 2014) or when modeling extremes (see Coles, 2001; Stephenson and Tawn, 2004) we can consider a definition for a marginally interpretable model based on relating the conditional quantiles to the marginal quantiles. Further research is needed to understand the form of the adjustments that arise in these settings.

Although we discuss fitting marginally interpretable GLMMs using Bayesian techniques, these models are also compatible with frequentist techniques. Regardless of the method used, the key is to include the adjustment at the appropriate step in the algorithm. Further, for models with a logit link and normal random effects, any model-fitting technique, be it Bayesian or frequentist, could benefit from our more accurate approach to evaluating the logistic-normal integral.

Another area where the marginally interpretable GLMM can improve inference is in hypothesis testing. When comparing two group means in a GLMM with a nonlinear link function, testing whether there is a difference between the two groups is not necessarily the same as testing whether the group means differ because the different groups could require different adjustments to preserve the marginal means. In other words, failure to account for the impact of the nonlinear link could lead one to test the wrong hypotheses. The importance of the adjustment was demonstrated with the rat teratology data in Section 5.1. Using the conventional GLMM, testing for a nonzero treatment effect yielded a different result than testing for a difference in the expected survival rates. Using the marginally interpretable GLMM avoids such inconsistencies because it makes the appropriate adjustments in the presence of random effects.

Appendix

Proof of Proposition 1: The marginal likelihood for each $Y_i, i = 1, \dots, N$, is

$$\begin{aligned}
f_Y(Y_i) &= \int f_{Y|U}(Y_i|U_i = u) f_U(u) du = \int E(Y_i|U_i = u)^{Y_i} \{1 - E(Y_i|U_i = u)\}^{1-Y_i} f_U(u) du \\
&= \int h(\mathbf{x}_i^T \boldsymbol{\beta} + u + a_i)^{Y_i} \{1 - h(\mathbf{x}_i^T \boldsymbol{\beta} + u + a_i)\}^{1-Y_i} f_U(u) du \\
&= h(\mathbf{x}_i^T \boldsymbol{\beta})^{Y_i} \{1 - h(\mathbf{x}_i^T \boldsymbol{\beta})\}^{1-Y_i},
\end{aligned}$$

where the final equality holds because the model satisfies (6).

Proof of Proposition 2: Suppose for each $i = 1, \dots, N$ we have $\mathbf{U}_i \sim N_q(\mathbf{0}, \boldsymbol{\Sigma})$ and write $\mathbf{U}_i = \boldsymbol{\Sigma}^{\frac{1}{2}} \mathbf{Z}_i$, where $\boldsymbol{\Sigma}^{\frac{1}{2}}$ is a square root matrix for $\boldsymbol{\Sigma}$ and $\mathbf{Z}_i \sim N_q(\mathbf{0}, \mathbf{I}_q)$. We can define a random variable $V = \mathbf{d}^T \boldsymbol{\Sigma}^{\frac{1}{2}} \mathbf{Z}$ such that $V \sim N(0, \tau^2)$, where $\tau^2 = \mathbf{d}^T \boldsymbol{\Sigma} \mathbf{d}$. It is also possible to define $q - 1$ additional random variables $\mathbf{W} = (W_1, \dots, W_{q-1})^T$ that span the orthogonal complement of V relative to \mathbb{R}^q such that \mathbf{W} follows a $(q - 1)$ -dimensional Normal distribution with density $f_{\mathbf{W}}(\cdot)$. Given such a V and \mathbf{W} , we have

$$\begin{aligned}
&\int_{\mathbb{R}^q} h(\kappa + \mathbf{d}^T \mathbf{u} + a) \left(\frac{1}{2\pi}\right)^{\frac{q}{2}} |\boldsymbol{\Sigma}|^{-1/2} \exp\left(-\frac{1}{2} \mathbf{u}^T \boldsymbol{\Sigma}^{-1} \mathbf{u}\right) d\mathbf{u} \\
&= \int_{\mathbb{R}^q} h(\kappa + \mathbf{d}^T \boldsymbol{\Sigma}^{1/2} \mathbf{z} + a) \left(\frac{1}{2\pi}\right)^{\frac{q}{2}} \exp\left(-\frac{1}{2} \mathbf{z}^T \mathbf{z}\right) d\mathbf{z}. \\
&= \int_{\mathbb{R}^q} h(\kappa + v + a) \frac{1}{\sqrt{2\pi\tau^2}} \exp\left(-\frac{1}{2\tau^2} v^2\right) f_{\mathbf{W}}(\mathbf{w}) d\mathbf{w} dv \\
&= \int_{\mathbb{R}} h(\kappa + v + a) \frac{1}{\sqrt{2\pi\tau^2}} \exp\left(-\frac{1}{2\tau^2} v^2\right) dv \int_{\mathbb{R}^{q-1}} f_{\mathbf{W}}(\mathbf{w}) d\mathbf{w} \\
&= \int_{\mathbb{R}} h(\kappa + v + a) \frac{1}{\sqrt{2\pi\tau^2}} \exp\left(-\frac{1}{2\tau^2} v^2\right) dv.
\end{aligned}$$

Proof of Proposition 3: For each $i = 1, \dots, N$, dividing both sides of (9) by $\exp(\mathbf{x}_i^T \boldsymbol{\beta} + \mathbf{d}_i^T \mathbf{a}_i)$ and then taking the natural logarithm of both sides yields

$$-\mathbf{d}_i^T \mathbf{a}_i = \log \left\{ \int \exp(\mathbf{d}_i^T \mathbf{u}) f_{\mathbf{U}}(\mathbf{u}) d\mathbf{u} \right\}.$$

Thus, $\mathbf{d}_i^T \mathbf{a}_i = -\log\{M_{\mathbf{U}}(\mathbf{d}_i)\}$, as required.

Proof of Corollary 1: This result follows immediately from Proposition 3. For each $i = 1, \dots, N$, if $\mathbf{d}_i^T \mathbf{a}_i$ exists, then $\mathbf{d}_i^T \mathbf{a}_i = -\log\{M_{\mathbf{U}}(\mathbf{d}_i)\}$, so $M_{\mathbf{U}}(\mathbf{d}_i)$ must exist. Conversely, if $M_{\mathbf{U}}(\mathbf{d}_i)$ exists, then $\mathbf{d}_i^T \mathbf{a}_i = -\log\{M_{\mathbf{U}}(\mathbf{d}_i)\}$ also exists because $M_{\mathbf{U}}(\mathbf{d}_i)$ is strictly positive.

Proof of Corollary 2: For this model, $d_i = 1$ for all $i = 1, \dots, N$ and $M_U(t) = \exp(\sigma^2 t^2 / 2)$. Thus, for all $i = 1, \dots, N$, $a_i = -\log\{M_U(d_i)\} = -\log\{\exp(\sigma^2 / 2)\} = -\sigma^2 / 2$, as required.

Proof of Theorem 1: To simplify notation, we suppress the subscript i . For any choice of $f_{\mathbf{U}}$,

$$-\infty < \ell = \ell \int f_{\mathbf{U}}(\mathbf{u}) d\mathbf{u} \leq \int h(\mathbf{x}^T \boldsymbol{\beta} + \mathbf{d}^T \mathbf{u} + \mathbf{d}^T \mathbf{a}) f_{\mathbf{U}}(\mathbf{u}) d\mathbf{u} \leq u \int f_{\mathbf{U}}(\mathbf{u}) d\mathbf{u} = u < \infty,$$

so the integral $\int h(\mathbf{x}^T \boldsymbol{\beta} + \mathbf{d}^T \mathbf{u} + \mathbf{d}^T \mathbf{a}) f_{\mathbf{U}}(\mathbf{u}) d\mathbf{u}$ exists. Further, $\int h(\mathbf{x}^T \boldsymbol{\beta} + \mathbf{d}^T \mathbf{u} + \mathbf{d}^T \mathbf{a}) f_{\mathbf{U}}(\mathbf{u}) d\mathbf{u}$ is a continuous function of $\mathbf{d}^T \mathbf{a}$ and, provided $\mathbf{d} \neq \mathbf{0}$, the following two limits hold:

$$\lim_{\mathbf{d}^T \mathbf{a} \rightarrow -\infty} \int h(\mathbf{x}^T \boldsymbol{\beta} + \mathbf{d}^T \mathbf{u} + \mathbf{d}^T \mathbf{a}) f_{\mathbf{U}}(\mathbf{u}) d\mathbf{u} = \ell,$$

$$\lim_{\mathbf{d}^T \mathbf{a} \rightarrow \infty} \int h(\mathbf{x}^T \boldsymbol{\beta} + \mathbf{d}^T \mathbf{u} + \mathbf{d}^T \mathbf{a}) f_{\mathbf{U}}(\mathbf{u}) d\mathbf{u} = u.$$

Thus, since $\ell \leq h(\mathbf{x}^T \boldsymbol{\beta}) \leq u$, for any value of $h(\mathbf{x}^T \boldsymbol{\beta})$ there exists an adjustment $\mathbf{d}^T \mathbf{a}$ such that (6) holds. For the case when $\mathbf{d} = \mathbf{0}$, $h(\mathbf{x}^T \boldsymbol{\beta} + \mathbf{d}^T \mathbf{u} + \mathbf{d}^T \mathbf{a}) = h(\mathbf{x}^T \boldsymbol{\beta})$ and (6) trivially holds.

Proof of Proposition 4: Let $\mathbf{U}_i \sim N_q(\mathbf{0}, \Sigma)$ and $\epsilon \sim N(0, 1)$, and define $W = \epsilon - \mathbf{d}_i^T \mathbf{U}_i$ so that $W \sim N(0, 1 + \mathbf{d}_i^T \Sigma \mathbf{d}_i)$. Then,

$$\begin{aligned} \int \Phi(\mathbf{x}_i^T \boldsymbol{\beta} + \mathbf{d}_i^T \mathbf{u} + \mathbf{d}_i^T \mathbf{a}_i) f_{\mathbf{U}}(\mathbf{u}) d\mathbf{u} &= \int P(\epsilon \leq \mathbf{x}_i^T \boldsymbol{\beta} + \mathbf{d}_i^T \mathbf{u} + \mathbf{d}_i^T \mathbf{a}_i) f_{\mathbf{U}}(\mathbf{u}) d\mathbf{u} \\ &= P(\epsilon \leq \mathbf{x}_i^T \boldsymbol{\beta} + \mathbf{d}_i^T \mathbf{U}_i + \mathbf{d}_i^T \mathbf{a}_i) = P(\epsilon - \mathbf{d}_i^T \mathbf{U}_i \leq \mathbf{x}_i^T \boldsymbol{\beta} + \mathbf{d}_i^T \mathbf{a}_i) \\ &= P(W \leq \mathbf{x}_i^T \boldsymbol{\beta} + \mathbf{d}_i^T \mathbf{a}_i) = \Phi \left\{ \frac{\mathbf{x}_i^T \boldsymbol{\beta} + \mathbf{d}_i^T \mathbf{a}_i}{(1 + \mathbf{d}_i^T \Sigma \mathbf{d}_i)^{1/2}} \right\}. \end{aligned}$$

Consequently,

$$\Phi(\mathbf{x}_i^T \boldsymbol{\beta}) = \int \Phi(\mathbf{x}_i^T \boldsymbol{\beta} + \mathbf{d}_i^T \mathbf{u} + \mathbf{d}_i^T \mathbf{a}_i) f_{\mathbf{U}}(\mathbf{u}) d\mathbf{u} = \Phi \left\{ \frac{\mathbf{x}_i^T \boldsymbol{\beta} + \mathbf{d}_i^T \mathbf{a}_i}{(1 + \mathbf{d}_i^T \Sigma \mathbf{d}_i)^{1/2}} \right\},$$

so $\mathbf{x}_i^T \boldsymbol{\beta} (1 + \mathbf{d}_i^T \Sigma \mathbf{d}_i)^{1/2} = \mathbf{x}_i^T \boldsymbol{\beta} + \mathbf{d}_i^T \mathbf{a}_i$, and $\mathbf{d}_i^T \mathbf{a}_i = \{(1 + \mathbf{d}_i^T \Sigma \mathbf{d}_i)^{1/2} - 1\} \mathbf{x}_i^T \boldsymbol{\beta}$, as required.

Proof of Proposition 5: We show the limit for $\kappa \rightarrow -\infty$. (The limit for $\kappa \rightarrow \infty$ follows from symmetry.) Let $\kappa = \mathbf{x}_i^T \boldsymbol{\beta}$, $a = \mathbf{d}_i^T \mathbf{a}_i$, and $\tau^2 = \mathbf{d}_i^T \Sigma \mathbf{d}_i$. Then from Proposition 2 it follows that,

$$\frac{e^\kappa}{1 + e^\kappa} = \int_{\mathbb{R}^q} \frac{e^{\kappa + \mathbf{u} + a}}{1 + e^{\kappa + \mathbf{u} + a}} f_{\mathbf{U}}(\mathbf{u}) d\mathbf{u} = \int_{\mathbb{R}} \frac{e^{\kappa + v + a}}{1 + e^{\kappa + v + a}} \frac{1}{\sqrt{2\pi\tau^2}} \exp\left(-\frac{1}{2\tau^2} v^2\right) dv$$

Dividing by $\exp(\kappa)$ on both sides and taking the limit as $\kappa \rightarrow -\infty$ we obtain

$$1 = e^a \int_{\mathbb{R}} e^v \frac{1}{\sqrt{2\pi\tau^2}} \exp\left(-\frac{1}{2\tau^2} v^2\right) dv,$$

which implies that $a = -\tau^2/2$. Thus, $\mathbf{d}_i^T \mathbf{a}_i = -\frac{1}{2} \mathbf{d}_i^T \Sigma \mathbf{d}_i$ for each $i = 1, \dots, N$, as required.

Proof of Proposition 6: Expanding the square in (16) and then rearranging terms yields

$$(\mathbf{d}_i^T \mathbf{a}_i)^2 + 2(\mathbf{x}_i^T \boldsymbol{\beta})(\mathbf{d}_i^T \mathbf{a}_i) + \text{Var}(\mathbf{d}_i^T \mathbf{U}_i) = 0.$$

Application of the quadratic formula then leads to

$$\mathbf{d}_i^T \mathbf{a}_i = \frac{1}{2}[-2\mathbf{x}_i^T \boldsymbol{\beta} \pm \{(2\mathbf{x}_i^T \boldsymbol{\beta})^2 - 4\text{Var}(\mathbf{d}_i^T \mathbf{U}_i)\}^{1/2}] = -\mathbf{x}_i^T \boldsymbol{\beta} \pm \{(\mathbf{x}_i^T \boldsymbol{\beta})^2 - \text{Var}(\mathbf{d}_i^T \mathbf{U}_i)\}^{1/2}.$$

Thus, subject to the constraint $\mathbf{x}_i^T \boldsymbol{\beta} + \mathbf{d}_i^T \mathbf{a}_i \geq 0$, we have

$$\mathbf{d}_i^T \mathbf{a}_i = -\mathbf{x}_i^T \boldsymbol{\beta} + \{(\mathbf{x}_i^T \boldsymbol{\beta})^2 - \text{Var}(\mathbf{d}_i^T \mathbf{U}_i)\}^{1/2}.$$

References

- Breslow, N. E. and Clayton, D. G. (1993). Approximate inference in generalized linear mixed models. *Journal of the American Statistical Association*, 88:9–25.
- Caffo, B., An, M.-W., and Rohde, C. (2007). Flexible random intercept models for binary outcomes using mixture of normals. *Computational Statistics and Data Analysis*, 51:5220–5235.
- Coles, S. (2001). *An Introduction to Statistical Modeling of Extreme Values*. Springer, New York, NY.
- Crouch, E. A. C. and Spiegelman, D. (1990). The evaluation of integrals of the form $\int_{-\infty}^{\infty} f(t) \exp(-t^2) dt$: Application to logistic-normal models. *Journal of the American Statistical Association*, 85:464–469.
- Dickey, J. M. (1971). The weighted likelihood ratio, linear hypotheses on normal location parameters. *The Annals of Mathematical Statistics*, 42:204–223.
- Diggle, P. J., Heagerty, P., Liang, K.-Y., and Zeger, S. L. (2002). *Analysis of Longitudinal Data*. Oxford University Press, Oxford, 2nd edition.

- Eddelbuettel, D. (2013). *Seamless R and C++ Integration with Rcpp*. Springer, New York, NY.
- Eddelbuettel, D. and François, R. (2011). Rcpp: Seamless R and C++ integration. *Journal of Statistical Software*, 40:1–18.
- Gamerman, D. (1997). Sampling from the posterior distribution in generalized linear mixed models. *Statistics and Computing*, 7:57–68.
- Geraci, M. and Bottai, M. (2014). Linear quantile mixed models. *Statistics and Computing*, 24:461–479.
- Gilmour, A. R., Anderson, R. D., and Rae, A. L. (1985). The analysis of binomial data by a generalized linear mixed model. *Biometrika*, 72:593–599.
- Heagerty, P. J. (1999). Marginally specified logistic-normal models for longitudinal binary data. *Biometrics*, 55:688–698.
- Heagerty, P. J. and Kurland, B. F. (2001). Misspecified maximum likelihood estimates and generalised linear mixed models. *Biometrika*, 88:973–985.
- Heagerty, P. J. and Zeger, S. L. (2000). Marginalized multilevel models and likelihood inference. *Statistical Science*, 15:1–19.
- Henderson, C. R. (1975). Best linear unbiased estimation and prediction under a selection model. *Biometrics*, 31:423–447.
- Henderson, C. R., Kempthorne, O., Searle, S. R., and von Krosigk, C. M. (1959). The estimation of environmental and genetic trends from records subject to culling. *Biometrics*, 15:192–218.
- Koenker, R. (2004). Quantile regression for longitudinal data. *Journal of Multivariate Analysis*, 91:74–89.
- Laird, N. M. and Ware, J. H. (1982). Random-effects models for longitudinal data. *Biometrics*, 38:963–974.

- Leppik, I., Dreifuss, F., Porter, R., Bowman, T., Santilli, N., Jacobs, M., Crosby, C., Cloyd, J., Stackman, J., Graves, N., Sutula, T., Welty, T., Vickery, J., Brundage, R., Gates, J., Gumnit, R., and Gutierrez, A. (1987). A controlled study of progabide in partial seizures: Methodology and results. *Neurology*, 37:963–968.
- Liang, K.-Y. and Hanfelt, J. (1994). On the use of the quasi-likelihood method in teratological experiments. *Biometrics*, 50:872–880.
- Liang, K.-Y. and Zeger, S. L. (1986). Longitudinal data analysis using generalized linear models. *Biometrika*, 73:13–22.
- Magder, L. S. and Zeger, S. L. (1996). A smooth nonparametric estimate of a mixing distribution using mixtures of Gaussians. *Journal of the American Statistical Association*, 91:1141–1151.
- McCullagh, P. and Nelder, J. A. (1989). *Generalized Linear Models*. Chapman & Hall/CRC, Boca Raton, FL, 2nd edition.
- McCulloch, C. E., Searle, S. R., and Neuhaus, J. M. (2008). *Generalized, Linear, and Mixed Models*. Wiley, Hoboken, NJ, 2nd edition.
- Miglioretti, D. L. and Heagerty, P. J. (2004). Marginal modeling of multilevel binary data with time-varying covariates. *Biostatistics*, 5:381–398.
- Monahan, J. F. and Stefanski, L. A. (1992). Normal scale mixture approximations to $f^*(z)$ and computation of the logistic-normal integral. In Balakrishnan, N., editor, *Handbook of the Logistic Distribution*, pages 529–540. Marcel Dekker, New York, NY.
- Nelder, J. A. and Wedderburn, R. W. M. (1972). Generalized linear models. *Journal of the Royal Statistical Society: Series A*, 135:370–384.
- Neuhaus, J. M. and Jewell, N. P. (1993). A geometric approach to assess bias due to omitted covariates in generalized linear models. *Biometrika*, 80:807–815.

- Neuhaus, J. M., Kalbfleisch, J. D., and Hauck, W. W. (1991). A comparison of cluster-specific and population-averaged approaches for analyzing correlated binary data. *International Statistical Review*, 59:25–35.
- Pirjol, D. (2013). The logistic-normal integral and its generalizations. *Journal of Computational and Applied Mathematics*, 237:460–469.
- R Core Team (2016). *R: A Language and Environment for Statistical Computing*. R Foundation for Statistical Computing, Vienna, Austria.
- Ritz, J. and Spiegelman, D. (2004). Equivalence of conditional and marginal regression models for clustered and longitudinal data. *Statistical Methods in Medical Research*, 13:309–323.
- Stephenson, A. and Tawn, J. (2004). Bayesian inference for extremes: accounting for the three extremal types. *Extremes*, 7:291–307.
- Stiratelli, R., Laird, N., and Ware, J. H. (1984). Random-effects models for serial observations with binary response. *Biometrics*, 40:961–971.
- Thall, P. F. and Vail, S. C. (1990). Some covariance models for longitudinal count data with overdispersion. *Biometrics*, 46:657–671.
- Verdinelli, I. and Wasserman, L. (1995). Computing Bayes factors using a generalization of the Savage-Dickey density ratio. *Journal of the American Statistical Association*, 90:614–618.
- Wang, Z. and Louis, T. A. (2003). Matching conditional and marginal shapes in binary random intercept models using a bridge distribution function. *Biometrika*, 90:765–775.
- Wang, Z. and Louis, T. A. (2004). Marginalized binary mixed-effects models with covariate-dependent random effects and likelihood inference. *Biometrics*, 60:884–891.

- Weil, C. S. (1970). Selection of the valid number of sampling units and a consideration of their combination in toxicological studies involving reproduction, teratogenesis or carcinogenesis. *Food and Cosmetics Toxicology*, 8:177–182.
- Zeger, S. L. and Karim, M. R. (1991). Generalized linear models with random effects; a Gibbs sampling approach. *Journal of the American Statistical Association*, 86:79–86.
- Zeger, S. L. and Liang, K.-Y. (1986). Longitudinal data analysis for discrete and continuous outcomes. *Biometrics*, 42:121–130.
- Zeger, S. L., Liang, K.-Y., and Albert, P. S. (1988). Models for longitudinal data: A generalized estimating equation approach. *Biometrics*, 44:1049–1060.